


Sparse Coding

Joan Bruna

(slides from Joan Bruna, Yann LeCun, Julien Mairal, Koray Kavukcuoglu)

Sparse Modeling


- ❖ Goal: Discover and exploit *patterns* in high-dimensional data:



$x_i \in \mathbb{R}^N$, N large

Sparse Modeling

- ❖ Goal: Discover and exploit *patterns* in high-dimensional data:



Sparsity in Signal Processing


$x \in \mathbb{R}^N$, $\mathcal{B}_N = \{b_1, \dots, b_N\}$ Basis of \mathbb{R}^N .

$$x = \sum_{i=1}^N \alpha_i b_i , \text{ i.e. } x = \mathcal{B}\alpha .$$


What is a “good” basis?



Pixels (Dirac basis)



Fourier




Wavelets

Sparsity in Signal Processing

More zeros with a redundant *dictionary*: $\mathcal{D} \in \mathbb{R}^{N \times M}$, $M > N$:

$$x \approx \mathcal{D}\alpha, \text{ with } \#\{|\alpha_i| > 0\} \text{ small}$$

Even more zeros with *adapted* dictionaries:



handwritten digit dictionary


Sparse Modeling

- ❖ Inverse Problems in Image processing
 - Denoising, Inpainting, Super-Resolution,...
- ❖ Unsupervised Learning
 - Video-surveillance, Source Separation,...
- ❖ Supervised Learning
 - Object Recognition

K-Means




K-Means




K-Means on MNIST

Y^i



$M^k \quad (K = 25)$



K-Means Algorithm

❖ Minimize wrt M :

$$E(M) = \sum_{i=1}^P \min_{k=1\dots,K} \|Y^i - M^k\|^2$$

❖ Lloyd Algorithm:

Fix M and update assignment:

For each $i \leq P$, $k(i) = \arg \min_k \|Y^i - M^k\|$.

Fix assignment and update M :

For each $k \leq K$, $S_k = \{i; \text{st } k(i) = k\}$.

$$M^k = \frac{1}{|S_k|} \sum_{i \in S_k} Y^i$$

How to use K-Means?

- * For a new sample Y , find k such that

$$k = \arg \min_k ||M^k - Y||_2^2$$

For $k = 1..K$

$$z_k = (M^k - Y)^T (M^k - Y)$$

End

$k = \text{index of } \min(z)$

- * Representation (1 of K)

$$z = [0 \ 0 \ 0 \ 0 \ \dots \ . \ z_k \ \dots \ \dots \ 0 \ 0 \ 0]$$

Sparse Coding

- * Represent an input vector using an **overcomplete** dictionary

$$\begin{array}{ccc}
 \begin{pmatrix} \cdot \\ \cdot \\ y_i \\ \cdot \\ \cdot \end{pmatrix} & \approx & \begin{pmatrix} \cdot & \cdot & \cdot & D_0^j & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & D_i^j & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{pmatrix} \begin{pmatrix} \cdot \\ \cdot \\ \cdot \\ z_j \\ \cdot \\ \cdot \\ \cdot \\ \cdot \end{pmatrix} \\
 \mathbf{Y} & & \mathbf{D} \\
 \text{Input} & & \text{Dictionary}
 \end{array}$$

$\# \text{ of dictionary elements} > \text{size of input}$
 $\# \text{ of zero elements} \gg \# \text{ of non-zero}$
 \mathbf{z} Representation (**sparse**)

- * Each \mathbf{Y} is represented using a linear combination of columns of \mathbf{D}
- * How do we calculate \mathbf{z} for a given \mathbf{Y} ?
- * How do we learn \mathbf{D} ?

Sparse Coding - L0

- 1) Find the sparsest solution that satisfies a given reconstruction error

$$\min ||z||_0 \quad s.t. \quad ||Y - \sum_i D^i z_i||_2^2 \leq \epsilon$$

- 2) Find the best k-sparse representation that minimizes reconstruction error

$$\min ||Y - \sum_i D^i z_i||_2^2 \quad s.t. \quad ||z||_0 = k$$

* L0 minimization requires search


* not tractable

Sparse Coding - L0


- * Matching Pursuit Algorithms offer greedy solution [Mallat and Zhang '93]
- * Greedily pick the dictionary element that reduces residual most
- * very fast, but unstable

```
Function MP (Y,D,n)
  R=Y,z=0
  for k=1..n
    i    = argmax(DTR)
    z_i  = DiTR
    R    = R - z_i Di
  end
```

Matching Pursuit Example




Y



D

Matching Pursuit Example




$k = 1$

Y

$Y - R$


R




Matching Pursuit Example




$k = 11$



Matching Pursuit Example





$k = 21$



Matching Pursuit Example

$k = 31$



Sparse Coding - L1

- * Relax L0 into closest convex penalty
- * Equality of minimum for L0 and L1 is proven under certain conditions [Donoho and Elad '03]

$$\frac{1}{2} || \underset{\substack{\uparrow \\ \text{Input}}}{Y} - \sum_i \underset{\substack{\uparrow \\ \text{Dictionary}}}{D^i} z_i ||_2^2 + \lambda \sum_i \underset{\substack{\uparrow \\ \text{Representation}}}{|z_i|}_1$$

- * Convex in **z** and **D** separately, not both
- * Fast algorithms exist for solving wrt **z**

Sparse Coding - LI

❖ Proximal Splitting Methods (gradient descent):

- 1st order methods
- Formulation for general functions of the form

$$\min_z F(z) + G(z)$$

F convex, smooth, with Lipschitz continuous gradient L_F .

G convex, continuous, and possibly non smooth.

- Quadratic Approximation of $F(z)+G(z)$ at z' :

$$Q(z, z') = F(z') + \langle z - z', \nabla F(z') \rangle + \frac{L}{2} \|z - z'\|^2 + G(z)$$

If $L \geq L_F$, then $\forall z, z', F(z) + G(z) \leq Q(z, z')$

Sparse Coding - LI

❖ Minimize $Q(z, z')$ with respect to z :

- Unique minimizer

- PROXIMAL OPERATOR of $F+G$:

$$\begin{aligned} P_L(z') &= \arg \min_z \{Q_L(z, z')\} \\ &= \arg \min_z G(z) + \frac{L}{2} \|z - (z' - \frac{1}{L} \nabla F(z'))\|^2 \end{aligned}$$

- Iterative Shrinkage Thresholding Algorithm (ISTA):

$$z^{k+1} = P_L(z^k)$$

$$z^{k+1} = \arg \min_z G(z) + \frac{L}{2} \|z - (z^k - \frac{1}{L} \nabla F(z^k))\|^2$$


Sparse Coding - LI

$$z^{k+1} \leftarrow \arg \min_z G(z) + \frac{L}{2} \left\| z - \left(z^k - \frac{1}{L} \nabla F(z^k) \right) \right\|^2$$

$\sum_i |z_i|_1$ \mathbf{t} $\frac{1}{2} \|Y - \sum_i D^i z_i\|_2^2$

$$z^{k+1} \leftarrow \arg \min_z \lambda \|z\|_1 + \frac{L}{2} \|z - t\|^2$$

$$z^{k+1} \leftarrow sh_{\lambda/L} \left(z^k - \frac{1}{L} \nabla F(z^k) \right)$$



$$z^{k+1} \leftarrow sh_{\lambda/L} \left(z^k - \frac{1}{L} D^T (Dz - Y) \right)$$

- * Loop until some convergence criterion is satisfied
- * How do we get L?

Sparse Coding - LI

- * **L** is the step size for gradient step
- * It is the smallest Lipschitz constant of the smooth function $F(z)$ and is equal to largest eigenvalue of $D^T D$
- * In practice, one does a line search

Function ISTA (Y, D)

$L > 0, c > 1, z = 0$

repeat

Search L s.t. $Q(z) > F(z) + G(z)$

$z^{k+1} = \text{sh}(z^k - 1/L D^T(Dz - Y))$

until convergence

Sparse Coding - LI

- ❖ Convergence Rate of ISTA:

$$\|z^k - z^*\| = O\left(\frac{1}{k}\right) .$$

- ❖ Can we find a faster first-order method?

FISTA

[Nesterov'83, Beck and Teboulle'09]

❖ Use Momentum in a clever way:

$$y^k = sh_{\lambda/L} \left(z^k - \frac{1}{L} D^T (D z^k - Y) \right)$$

$$t_{(k+1)} = \frac{1 + \sqrt{1 + 4t_{(k)}^2}}{2}$$

$$z^k = y^k + \frac{t_{(k)} - 1}{t_{(k+1)}} (y^k - y^{k-1})$$

❖ First Order, Optimal convergence for 1st order method:

$$\|z^k - z^*\| = O\left(\frac{1}{k^2}\right) .$$

Sparse Coding - Learning

- * How about **D**?

- * We want to learn it

- * Adapt to data

- * Use online learning for **D**

- * Per sample energy

$$E(Y, z, D) = \frac{1}{2} \|Y - \sum_i D^i z_i\|_2^2 + \lambda \sum_i |z_i|_1$$

- * Loss

$$L(Y, D) = \frac{1}{|\mathcal{Y}|} \sum_{Y \in \mathcal{Y}} E(Y, z, D)$$

Sparse Coding - Learning

* For each sample, $Y \in \mathcal{Y}$

1. do inference

minimize $\mathbf{E}(\mathbf{Y}, \mathbf{z}, \mathbf{D})$ wrt \mathbf{z} (use any SC algo)

2. update parameters


$$D \leftarrow D - \eta \frac{\partial E}{\partial D}$$

3. Constrain elements of D to be unit norm

* dictionary elements grow, z gets smaller,
sparsity term gets discarded


Sparse Coding - Learning

- * Learned dictionary **D** by training in natural image patches



Sparse Coding - Learning

- * Learned dictionary D by training on MNIST digits



Sparse Coding


- * Cool, available software?
 - * <http://www.di.ens.fr/willow/SPAMS/>
 - * <http://cs.nyu.edu/~koray>
- * Applications
 - * Image denoising
 - * Inpainting
 - * Classification
 - * Recognition
 - * ...

Image Processing Applications

- * Slides from Julien Mairal
- * <http://www.di.ens.fr/~mairal/resources/pdf/ERMITES10.pdf>

Sparse representations for image restoration

[Mairal, Sapiro, and Elad, 2008d]



Sparse representations for image restoration

Inpainting, [Mairal, Elad, and Sapiro, 2008b]



Since 1699, when French explorers landed at the great bend of the Mississippi River and celebrated the first Mardi Gras in North America, New Orleans has brewed a fascinating melange of cultures. It was French, then Spanish, then French again, then sold to the United States. Through all these years, and even into the 1900s, others arrived from everywhere: Acadians (Cajuns), Africans, indige-

Sparse representations for image restoration

Inpainting, [Mairal, Elad, and Sapiro, 2008b]



Sparse representations for video restoration

Key ideas for video processing

[Protter and Elad, 2009]

- Using a 3D dictionary.
- Processing of many frames at the same time.
- Dictionary propagation.

Sparse representations for image restoration

Inpainting, [Mairal, Sapiro, and Elad, 2008d]




Figure: Inpainting results.

Sparse representations for image restoration

Inpainting, [Mairal, Sapiro, and Elad, 2008d]




Figure: Inpainting results.

Sparse representations for image restoration

Inpainting, [Mairal, Sapiro, and Elad, 2008d]




Figure: Inpainting results.

Sparse representations for image restoration

Inpainting, [Mairal, Sapiro, and Elad, 2008d]




Figure: Inpainting results.

Sparse representations for image restoration

Inpainting, [Mairal, Sapiro, and Elad, 2008d]




Figure: Inpainting results.

Sparse representations for image restoration

Color video denoising, [Mairal, Sapiro, and Elad, 2008d]




Figure: Denoising results. $\sigma = 25$

Sparse representations for image restoration

Color video denoising, [Mairal, Sapiro, and Elad, 2008d]



Figure: Denoising results. $\sigma = 25$

Sparse representations for image restoration

Color video denoising, [Mairal, Sapiro, and Elad, 2008d]



Figure: Denoising results. $\sigma = 25$

Sparse representations for image restoration

Color video denoising, [Mairal, Sapiro, and Elad, 2008d]



Figure: Denoising results. $\sigma = 25$

Sparse representations for image restoration

Color video denoising, [Mairal, Sapiro, and Elad, 2008d]



Figure: Denoising results. $\sigma = 25$

Digital Zooming

Couzinie-Devy, 2010, Original



Digital Zooming

Couzinie-Devy, 2010, Bicubic




Digital Zooming

Couzinie-Devy, 2010, Proposed method




Digital Zooming

Couzinie-Devy, 2010, Original




Digital Zooming

Couzinie-Devy, 2010, Bicubic



Digital Zooming

Couzinie-Devy, 2010, Proposed approach



Inverse half-toning

Original



Inverse half-toning

Reconstructed image



Inverse half-toning

Original



Inverse half-toning

Reconstructed image



Inverse half-toning

Original




Inverse half-toning

Reconstructed image



Inverse half-toning

Original



Inverse half-toning

Reconstructed image




Inverse half-toning

Original



Inverse half-toning

Reconstructed image



Optimization for Dictionary Learning

Inpainting a 12-Mpixel photograph

THE SALINAS VALLEY is in Northern California. It is a long narrow swale between two ranges of mountains, and the Salinas River winds and twists up the center until it falls at last into Monterey Bay.

I remember my childhood games for grasses and secret flowers. I remember where a toad may live and what time the birds awaken in the summer and what trees and seasons smelled like-how people looked and walked and smelled even. The memory of odors is very rich.

I remember that the Gabilan Mountains to the east of the valley were light gay mountains full of sun and loveliness and a kind of invitation, so that you wanted to climb into their warm foothills almost as you want to climb into the lap of a beloved mother. They were beckoning mountains with a brown grass love. The Santa Lucias stood up against the sky to the west and kept the valley from the open sea, and they were dark and brooding unfriendly and dangerous. I always found in myself a dread of west and a love of east. Where I ever got such an idea I cannot say, unless it could be that the morning came over the peaks of the Gabilans and the night drifted back from the ridges of the Santa Lucias. It may be that the birth and death of the day had some part in my feeling about the two ranges of mountains.

From both sides of the valley little streams slipped out of the hill canyons and fell into the bed of the Salinas River. In the winter of wet years the streams ran full-freshet, and they swelled the river until sometimes it raged and boiled, bank full, and then it was a destroyer. The river tore the edges of the farm lands and washed whole acres down; it toppled barns and houses into itself, to go floating and bobbing away. It trapped cows and pigs and sheep and drowned them in its muddy brown water and carried them to the sea. Then when the late spring came, the river drew in from its edges and the sand banks appeared. And in the summer the river didn't run at all above ground. Some pools would be left in the deep swirl places under a high bank. The tules and grasses grew back, and willows straightened up with the flood debris in their upper branches. The Salinas was only a part-time river. The summer sun drove it underground. It was not a fast river at all, but it was the only one we had and so we boasted about it how dangerous it was in a wet winter and how dry it was in a dry summer. You can boast about anything if it's all you have. Maybe the less you have, the more you are required to boast.

The floor of the Salinas Valley, between the ranges and below the foothills, is level because this valley used to be the bottom of a hundred-mile inlet from the sea. The river mouth at Moss Landing was centuries ago the entrance to this long inland water. Once, fifty miles down the valley, my father bored a well. The drill came up first with topsoil and then with gravel and then with white sea sand full of shells and even pebbles.

Optimization for Dictionary Learning

Inpainting a 12-Mpixel photograph



Optimization for Dictionary Learning

Inpainting a 12-Mpixel photograph



Optimization for Dictionary Learning

Inpainting a 12-Mpixel photograph



Sparse Coding for Recognition

- * Recognition requires two basic operations
 - * Feature extraction
 - * Classification


- * Feature extraction

$$\mathcal{F}(Y) : Y \mapsto z$$

- * SIFT, HoG,
- * Convolutional Nets (w/o the last layer)
- * Use sparse coding inference as feature extractor
 - * MP(Y) or ISTA(Y)

Sparse Coding for Recognition


❖ Unsupervised Learning / Manifold Learning:



$$z^*(x) = \arg \min_z \frac{1}{2} \|x - Dz\|^2 + \lambda \|z\|_1 .$$

Sparse Coding for Recognition


❖ Unsupervised Learning / Manifold Learning:



$$z^*(x) = \arg \min_z \frac{1}{2} \|x - Dz\|^2 + \lambda \|z\|_1 .$$

Sparse Coding for Recognition

❖ Unsupervised Learning / Manifold Learning:




$z^*(x_1), z^*(x_2)$ same support

$$\|z^*(x_1) - z^*(x_2)\| \approx \|x_1 - x_2\|$$

$$z^*(x) = \arg \min_z \frac{1}{2} \|x - Dz\|^2 + \lambda \|z\|_1 .$$

Sparse Coding for Recognition

❖ Unsupervised Learning / Manifold Learning:




$z^*(x_1), z^*(x_2)$ different support

$$\|z^*(x_3) - z^*(x_4)\| \gg \|x_3 - x_4\|$$

$$z^*(x) = \arg \min_z \frac{1}{2} \|x - Dz\|^2 + \lambda \|z\|_1 .$$

Sparse Coding for Recognition

❖ Unsupervised Learning / Manifold Learning:



$z^*(x_1), z^*(x_2)$ different support

intrinsic metric $\approx \|z^*(x_3) - z^*(x_4)\| \gg \|x_3 - x_4\|$

$$z^*(x) = \arg \min_z \frac{1}{2} \|x - Dz\|^2 + \lambda \|z\|_1 .$$

Sparse Coding for Recognition

- * Mid-level feature extraction

$$\mathcal{F}(z) : z \mapsto z'$$

- * First layer

- * SIFT, HoG,

- * Second layer

- * Quantize into a visual wordbook

- * Do Sparse coding to learn wordbook

Learning Codebooks for Image Classification

Let an image be represented by a set of low-level descriptors \mathbf{y}_i at N locations identified with their indices $i = 1, \dots, N$.

- hard-quantization:

$$\mathbf{y}_i \approx \mathbf{D}\boldsymbol{\alpha}_i, \quad \alpha_i \in \{0, 1\}^p \quad \text{and} \quad \sum_{j=1}^p \alpha_i[j] = 1$$

- soft-quantization:

$$\alpha_i[j] = \frac{e^{-\beta \|\mathbf{y}_i - \mathbf{d}_j\|_2^2}}{\sum_{k=1}^p e^{-\beta \|\mathbf{y}_i - \mathbf{d}_k\|_2^2}}$$

- sparse coding:

$$\mathbf{y}_i \approx \mathbf{D}\boldsymbol{\alpha}_i, \quad \boldsymbol{\alpha}_i = \arg \min_{\boldsymbol{\alpha}} \frac{1}{2} \|\mathbf{y}_i - \mathbf{D}\boldsymbol{\alpha}\|_2^2 + \lambda \|\boldsymbol{\alpha}\|_1$$

Learning Codebooks for Image Classification

Table from Boureau et al. [2010]

Method	Caltech-101, 30 training examples		15 Scenes, 100 training examples	
	Average Pool	Max Pool	Average Pool	Max Pool
Results with basic features, SIFT extracted each 8 pixels				
Hard quantization, linear kernel	51.4 ± 0.9 [256]	64.3 ± 0.9 [256]	73.9 ± 0.9 [1024]	80.1 ± 0.6 [1024]
Hard quantization, intersection kernel	64.2 ± 1.0 [256] (1)	64.3 ± 0.9 [256]	80.8 ± 0.4 [256] (1)	80.1 ± 0.6 [1024]
Soft quantization, linear kernel	57.9 ± 1.5 [1024]	69.0 ± 0.8 [256]	75.6 ± 0.5 [1024]	81.4 ± 0.6 [1024]
Soft quantization, intersection kernel	66.1 ± 1.2 [512] (2)	70.6 ± 1.0 [1024]	81.2 ± 0.4 [1024] (2)	83.0 ± 0.7 [1024]
Sparse codes, linear kernel	61.3 ± 1.3 [1024]	71.5 ± 1.1 [1024] (3)	76.9 ± 0.6 [1024]	83.1 ± 0.6 [1024] (3)
Sparse codes, intersection kernel	70.3 ± 1.3 [1024]	71.8 ± 1.0 [1024] (4)	83.2 ± 0.4 [1024]	84.1 ± 0.5 [1024] (4)
Results with macrofeatures and denser SIFT sampling				
Hard quantization, linear kernel	55.6 ± 1.6 [256]	70.9 ± 1.0 [1024]	74.0 ± 0.5 [1024]	80.1 ± 0.5 [1024]
Hard quantization, intersection kernel	68.8 ± 1.4 [512]	70.9 ± 1.0 [1024]	81.0 ± 0.5 [1024]	80.1 ± 0.5 [1024]
Soft quantization, linear kernel	61.6 ± 1.6 [1024]	71.5 ± 1.0 [1024]	76.4 ± 0.7 [1024]	81.5 ± 0.4 [1024]
Soft quantization, intersection kernel	70.1 ± 1.3 [1024]	73.2 ± 1.0 [1024]	81.8 ± 0.4 [1024]	83.0 ± 0.4 [1024]
Sparse codes, linear kernel	65.7 ± 1.4 [1024]	75.1 ± 0.9 [1024]	78.2 ± 0.7 [1024]	83.6 ± 0.4 [1024]
Sparse codes, intersection kernel	73.7 ± 1.3 [1024]	75.7 ± 1.1 [1024]	83.5 ± 0.4 [1024]	84.3 ± 0.5 [1024]

	Unsup	Discr
Linear	83.6 ± 0.4	84.9 ± 0.3
Intersect	84.3 ± 0.5	84.7 ± 0.4

Yang et al. [2009] have won the PASCAL VOC'09 challenge using this kind of techniques.

Training Sparse Inference

- ❖ Sparse Coding solution requires per-datapoint optimization: L0 (very expensive) or L1 (expensive)
- ❖ Efficient algorithms, but still slow
- ❖ Can we train a feed-forward Neural Network that predicts sparse codes for feature extraction?
 - Predictive Sparse Decomposition (PSD) [Kavukcuoglu et al'08]
 - LISTA [Gregor, YLC'10]

Predictive Sparse Decomposition

$$E(Y, z, D, K) = \underbrace{\frac{1}{2} ||Y - \sum_i D^i z_i||_2^2}_{\text{Sparse Coding}} + \lambda \sum_i |z_i|_1 + \underbrace{\beta ||z - C(Y; K)||_2^2}_{\text{Prediction}}$$

Sparse Coding

Prediction


$$C(Y; K) = g \cdot \tanh(y * k)$$

* Learning


* For each sample from data, do:

1. Fix K and D , minimize to get optimal z
2. Using the optimal value of z update D and K
3. Scale elements of D to be unit norm.

Predictive Sparse Decomposition




Encoder (K)



Decoder (D)

- * 12 x 12 natural image patches
- * 256 dictionary elements



Recognition - CI01




- * **Optimal** (Feature Sign, Lee'07) vs **PSD** features
- * PSD features perform slightly better
- * Naturally optimal point of sparsity
- * After 64 features not much gain
- * PSD features are hundreds of times faster

Pedestrian Detection

- * Convolutional Predictive Sparse Decomposition
- * 2 layer architecture
- * INRIA Dataset
- * Unsupervised Learning improves by 20%




Pedestrian Detection



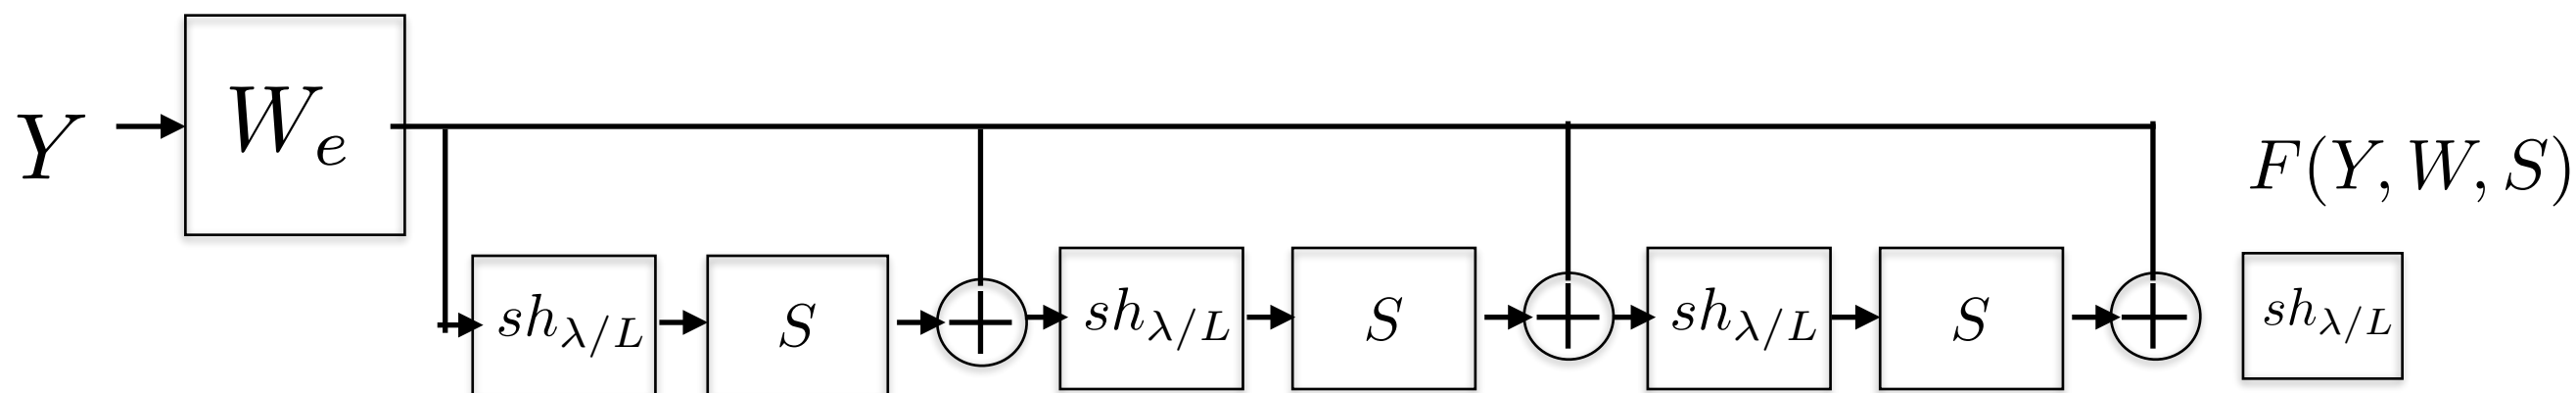
LISTA

[Gregor & LeCun'10]

❖ Proximal splitting algorithms (ISTA, FISTA):



❖ Idea: Unroll the recurrence K steps:



Trainable parameters: W_e , S , λ .

Training LISTA

$$E(Y, z, W, S) = \frac{1}{2} \sum_i \|Y_i - Dz_i\|^2 + \lambda \sum_i \|z_i\|_1 + \beta/2 \|z_i - F(Y_i, S, W)\|^2 .$$

❖ Results [from Gregor, YLC'10]:

Table 1. Prediction error (squared error between the optimal code and the predicted codes) for different non-linearities of the baseline encoder (7), and for LISTA and FISTA. LISTA produces a much better estimate after one iteration than FISTA.

NON-LINEARITY	100 UNITS	400 UNITS
$D \tanh(x)$	8.6	10.7
$D(\tanh(x + u) + \tanh(x - u))$	3.33	4.62
$h_\alpha(x)$	3.29	4.82
LISTA 1 ITERATION	1.50	2.45
LISTA 3 ITERATIONS	0.98	2.12
LISTA 7 ITERATIONS	0.52	1.62
FISTA 1 ITERATION	21.	22.

❖ Can perform “explaining away” unlike PSD with matrix S